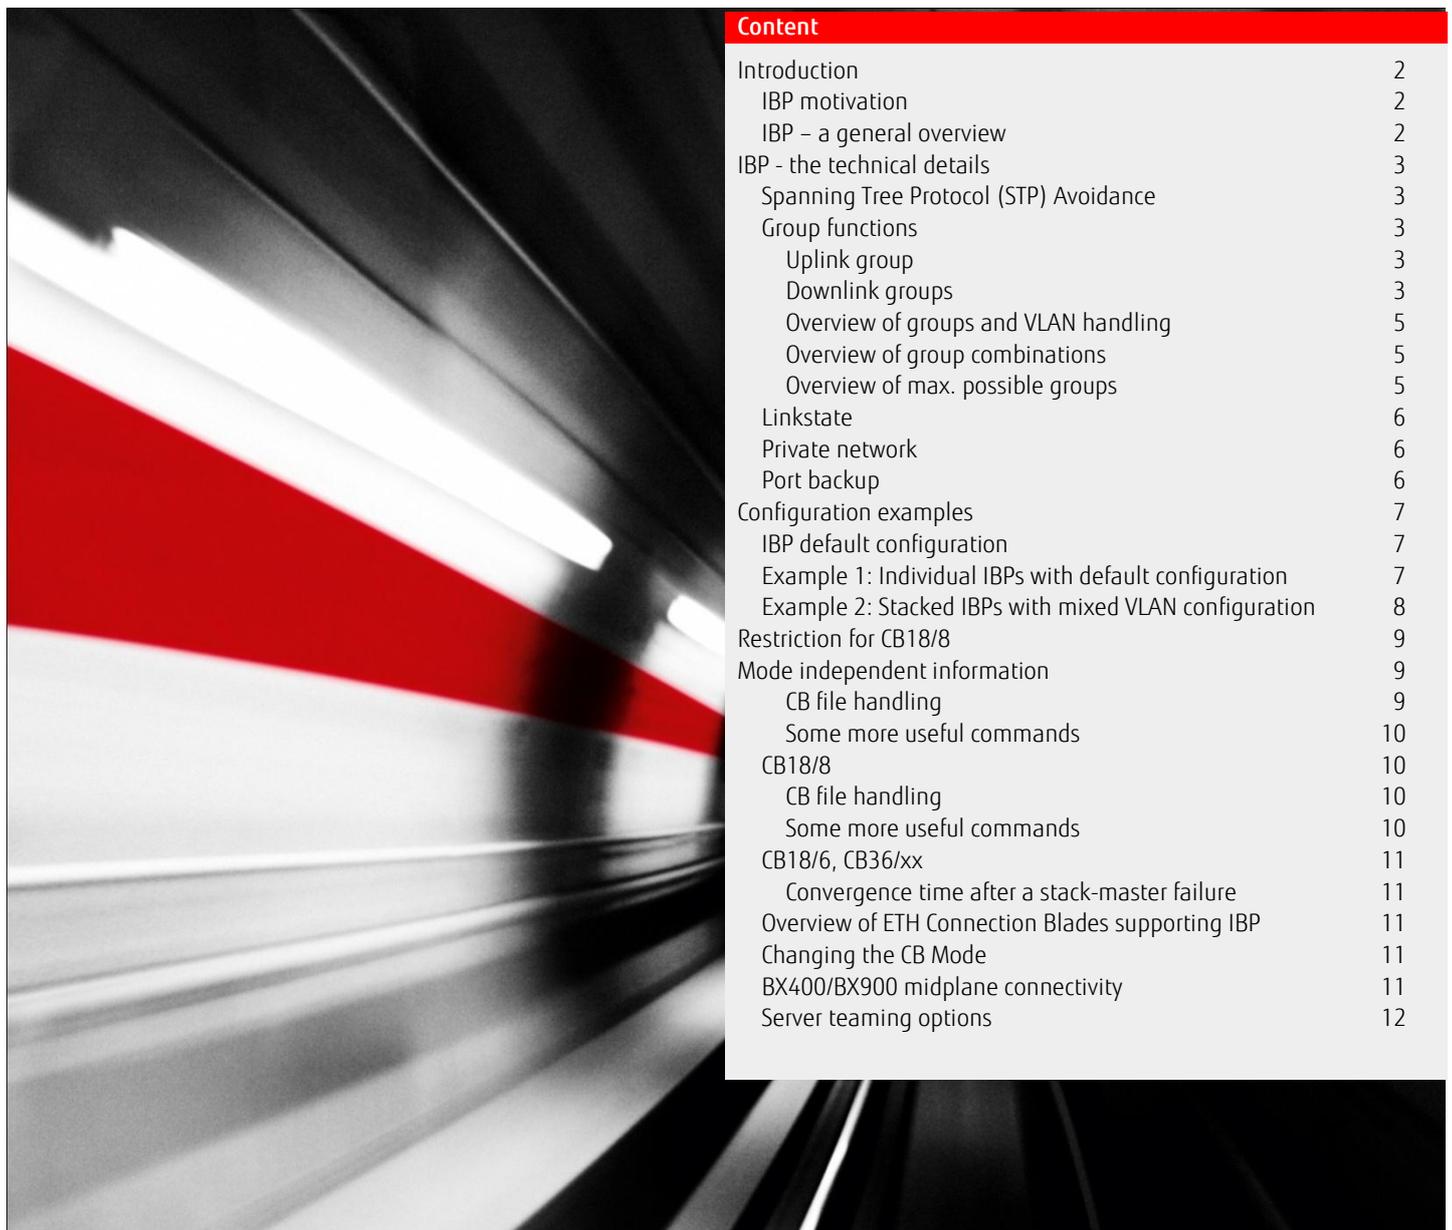


Technical White paper

INTELLIGENT BLADE PANEL (IBP)

for PRIMERGY BX ETH CB 18/8, 18/6, 36/12, 36/8+2

Blade servers in general are – even though based on industry standard components - systems with vendor-specific architectures, making them to a certain degree more particular than single rack or tower servers. This also applies to the network connectivity of such systems. Fujitsu now presents a simple solution.

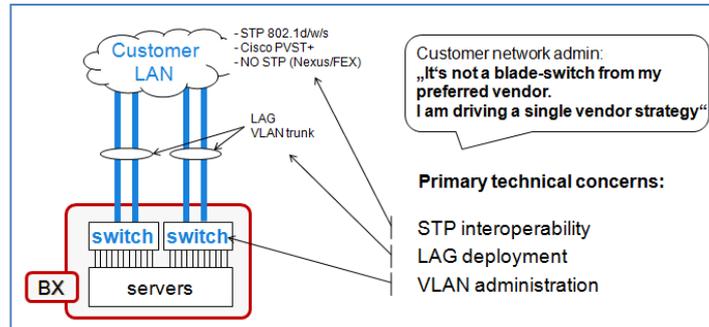


Content	
Introduction	2
IBP motivation	2
IBP – a general overview	2
IBP - the technical details	3
Spanning Tree Protocol (STP) Avoidance	3
Group functions	3
Uplink group	3
Downlink groups	3
Overview of groups and VLAN handling	5
Overview of group combinations	5
Overview of max. possible groups	5
Linkstate	6
Private network	6
Port backup	6
Configuration examples	7
IBP default configuration	7
Example 1: Individual IBPs with default configuration	7
Example 2: Stacked IBPs with mixed VLAN configuration	8
Restriction for CB18/8	9
Mode independent information	9
CB file handling	9
Some more useful commands	10
CB18/8	10
CB file handling	10
Some more useful commands	10
CB18/6, CB36/xx	11
Convergence time after a stack-master failure	11
Overview of ETH Connection Blades supporting IBP	11
Changing the CB Mode	11
BX400/BX900 midplane connectivity	11
Server teaming options	12

Introduction

IBP motivation

Network administrators prefer supplying their data centers with all identical devices and brands (e. g. all from Cisco). This is easy to understand: the administrators are well versed in "their" equipment, and inter-operability between the devices is ensured. A closer look at that "single-vendor strategy" means that the technical concerns are interoperability, deployment and administration of a non-branded switch. This is clearly to be seen with STP, LAG and VLAN, which are almost always in place. Although these technologies are standardized, the implementation as well as the configuration tasks may differ from switch to switch; it requires dedicated configuration tasks and these are not always compatible.



In order to support their customers, Fujitsu introduced two further operational modes for all PRIMERGY BX400/BX900 Ethernet switch blades (hereinafter referred to as connection blades = CB) in addition to the standard layer 2 switch mode of these components: These are the IBP (Intelligent Blade Panel) mode and the EHM (End-Host Mode). This very simple way of establishing connections and administering the CBs means that these two operational modes enable customers to easily and securely integrate PRIMERGY BX400/BX900 blade servers into their networks.

Note: This White Paper focuses on IBP -mode only.

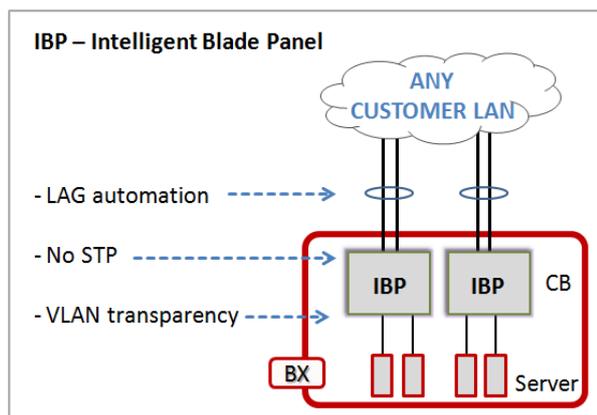
For an overview of all available modes EHM | IBP | Switch, see the White Paper "[PRIMERGY BX integration in any data center network](#)".

IBP – a general overview

IBP is based on common standard switch behavior with some elementary improvements. That is, IBP...

- does not interfere with the customer STP domain nor does it introduce data loops. It has no STP implemented and also avoids loops inherently with disabled switching between uplink ports.
- makes any VLAN administration needless. With the standard provider-feature 802.1q-in-q tagging, the CB provides transparent VLAN forwarding. With that, VLANs configured on the Customer LAN switch are automatically passed-thru to the blade servers respectively to the hypervisor virtual switch. As an option, VLANs can be assigned individually, like on a standard switch.
- does not burden additional effort for LAG deployment. LAG creation is greatly automated.
- has an advanced cli for easy and save feature configuration.
- is designed to be in the server administration domain due to its advanced cli.
- provides a very useful configuration initially, like LAG, VLAN transparency and redundancy.

These characteristics mean that the BX400/BX900 Ethernet Connection Blades with IBP-mode are very easy and secure for integration in any data center networking environment.



IBP - the technical details

Spanning Tree Protocol (STP) Avoidance

In IBP, STP is not implemented. No STP BPDU packets are sent by the IBP. This avoids STP interference with the customer network domain and supports proprietary implementations, such as Cisco PVST and Cisco Nexus FEX which has BPDU-guard hardcoded on its downlink ports. It is not possible to enable STP on the IBP. Data loops are avoided with disabled switching between IBP uplink ports.

Group functions

The IBP uses the so-called "Groups" functions to provide the option of flexibly allocating ports to each other. There is one uplink group called "uplink-set" and four downlink groups (port-group, VLAN-group, service-LAN, service-VLAN) described below in more detail.

Note

The uplink-set contains uplink ports; the ports of an uplink-set automatically form a LAG (static or dynamic). Downlink groups contain downlink ports; they differ in their VLAN behavior and are usually associated with an uplink-set.

Uplink group

The following applies for an uplink group:

- IBP uplink ports of the same uplink-set automatically form an LAG according to standard 802.3ad, statically or dynamically (LACP)
- IBP uplink ports without affiliation to an uplink-set have the link status "down"
- Ports of a joint uplink-set must be identically configured (speed / duplex / flow control)
- Up to eight ports can belong to an uplink-set
- An uplink-set can be connected with one or more downlink port groups

There is only one uplink group:

- uplink-set

Uplink-set creation and uplink port assignment can be carried out with the single command

#uplink-set <name>

Configuration example

Create uplink-set C-LAN:	<i>(config)#uplink-set C-LAN</i>
Add IBP uplink ports to uplink-set C-LAN:	<i>(Interface)#uplink-set C-LAN</i>

Downlink groups

The following applies for a downlink group:

- Downlink groups differ in their VLAN behavior. VLANs are according to standard 802.1q
- A downlink group can only be connected with a single uplink-set
- Ports of the same downlink group are always connected with each other, except in the "Private Network" configuration
- Ports of different groups are never connected with each other
- Downlink ports, which do not belong to a group, have the link status "down"
- Ports of a joint downlink group must be identically configured (speed / duplex / flow control)

There are four downlink groups (port-group, VLAN-group, service-LAN, service-VLAN) described below in more detail:

- Port-group

A port-group is VLAN independent and VLAN transparent. VLAN tagged and untagged packets are forwarded transparently. Port-group creation and uplink-set association can be carried out with the command:

#port-group <port-group name> <uplink-set name>

Downlink port assignment can be carried out with the command:

(interface)#port-group <port-group name>

Configuration example

Create uplink-set C-LAN:	<i>(config)#uplink-set C-LAN</i>
Add IBP uplink ports to uplink-set C-LAN:	<i>(Interface)#uplink-set C-LAN</i>
Create port group Server and connect with uplink-set C-LAN:	<i>(config)#port-group Server C-LAN</i>
Remove IBP downlink port from default port group:	<i>(Interface)#no port-group</i>
Add IBP downlink ports to port group Server:	<i>(Interface)#port-group Server</i>

- VLAN-group

A VLAN-group is allocated to a dedicated VLAN. Downlink ports of a VLAN-group can be untagged or tagged; uplink ports of the uplink-set associated with a VLAN-group are tagged, with native (untagged) option.

VLAN-group creation and uplink-set association can be carried out with the following command:

(config)#vlan-group <vlan-group name> <vlan-id> <uplink-set name>

Option: To configure native (untagged) VLAN for the uplink-set (default: vlan-group name with VLAN-id 1):

(config)#vlan-group-nativevlan <vlan-group name>

Downlink port assignment can be carried out with the command:

(interface)#vlan-group <vlan-group name> {tagged}

Configuration example

Create uplink-set C-LAN:	<i>(config)#uplink-set C-LAN</i>
Add IBP uplink ports to uplink-set C-LAN:	<i>(Interface)#uplink-set C-LAN</i>
Create VLAN group Server11 with VLAN-id 11 and connect with uplink-set C-LAN:	<i>(config)#vlan-group Server11 11 C-LAN</i>
Remove IBP downlink port from default port group:	<i>(Interface)#no port-group</i>
Add IBP downlink ports to VLAN group Server11:	<i>(Interface)#vlan-group Server11</i>

- Service-LAN

A service-LAN is allocated to a dedicated VLAN. Downlink ports of a service-LAN are tagged; uplink ports of the uplink-set associated with a service-LAN are untagged.

Service-LAN creation and uplink-set association can be carried out with the following command:

(config)#svc-lan <service-lan name> <vlan-id> <uplink-set name>

Downlink port assignment can be carried out with the command:

(interface)#svc-lan <service-lan name>

Configuration example

Create uplink-set C-LAN:	<i>(config)#uplink-set C-LAN</i>
Add IBP uplink ports to uplink-set C-LAN:	<i>(Interface)#uplink-set C-LAN</i>
Create service LAN Service99 with VLAN-id 99 and connect with uplink-set C-LAN:	<i>(config)#svc-lan Service99 99 C-LAN</i>
Remove IBP downlink port from default port group:	<i>(Interface)#no port-group</i>
Add IBP downlink ports to service LAN Service99:	<i>(Interface)#svc-lan Service99</i>

- Service-VLAN

A service-VLAN is allocated to a dedicated VLAN. Downlink ports of a service-VLAN are tagged; uplink ports of the uplink-set associated with a service-VLAN are tagged.

Service-VLAN creation and uplink-set association can be carried out with the following command:

(config)#svc-vlan <service-vlan name> <vlan-id> <uplink-set name>

Downlink port assignment can be carried out with the command:

(interface)#svc-vlan <service-vlan name>

Configuration example

Create uplink-set C-LAN:	<i>(config)#uplink-set C-LAN</i>
Add IBP uplink ports to uplink-set C-LAN:	<i>(Interface)#uplink-set C-LAN</i>
Create service VLAN Service333 with VLAN-ID 333 and connect with uplink-set C-LAN:	<i>(config)#svc-vlan Service333 333 C-LAN</i>
Remove IBP downlink port from default port group:	<i>(Interface)#no port-group</i>
Add IBP downlink ports to service VLAN Service333:	<i>(Interface)#svc-vlan Service333</i>

Overview of groups and VLAN handling

As described previously, the downlink groups differ in their VLAN behavior. The picture below gives an overview of all downlink groups and its VLAN behavior of the IBP downlink- and uplink ports.

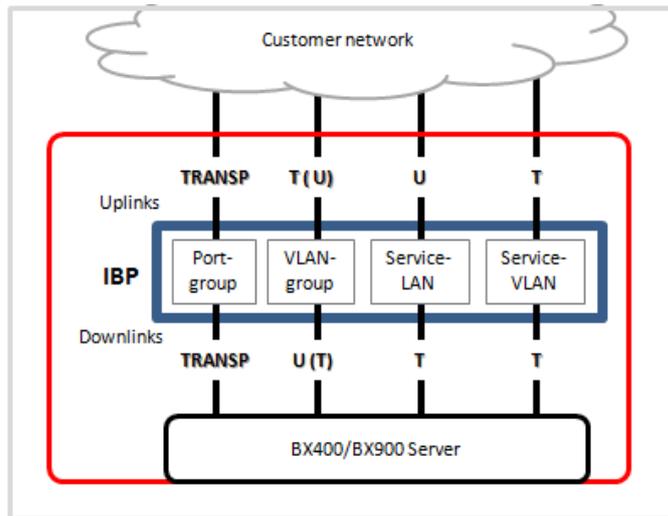
TRANSP: VLAN transparent; VLANs configured on the upstream (customer) switch are forwarded to the server (hypervisor Vswitch) transparently

T (U): Accepts / sends VLAN tagged packets only with the native (untagged) VLAN option

U (T): Accepts / sends VLAN untagged packets only with the option to tag VLANs

U: Accepts / sends VLAN untagged packets only

T: Accepts / sends VLAN tagged packets only



Overview of group combinations

Valid downlink group combinations

	Port-group	VLAN-group	Service-LAN	Service-VLAN
Port-group	no	no	yes	yes
VLAN-group	no	yes *	yes	yes
Service-LAN	yes	yes	yes	yes
Service-VLAN	yes	yes	yes	yes

* Requires version 6.xx or higher (CB36/xx) respectively v3.xx or higher (CB18/8)

Valid uplink group combinations

	Port-group	VLAN-group	Service-LAN	Service-VLAN
Port-group	no	no	no	yes
VLAN-group	no	yes	no	yes
Service-LAN	no	no	no	no
Service-VLAN	yes	yes	no	yes

Examples:

A downlink port can be member of a Port-group and a Service-LAN, but not of a Port-group and a VLAN-group at the same time.

An uplink-set can be associated with a Port-group and a Service-VLAN, but not with a Port-group and a VLAN-group at the same time.

Overview of max. possible groups

	CB 36/12 (SB11a)	CB 36/8+2 (SB11)	CB18/8 (SBAX2)
Max. uplink-set	12	10	8
Max. port-group	12	10	8
Max. VLAN-group	30	30	1000
Max. service-LAN	12	10	8
Max. service-VLAN	128	128	1000

Linkstate

Linkstate monitors the IBP uplink ports and thus the connection between the IBP and the customer switch. If a connection fails (all of the physical connections, or logical, i.e. LAG/LACP), Linkstate sets all downlink ports of those port groups to "link down" that are connected with the appropriate uplink-set. In this way, server teaming can react spontaneously and, if necessary, use a redundant connection. Linkstate is activated by default. It can be carried out with the command **(config)#linkstate <uplink-set name>**

Configuration example

Activate link state for uplink-set C-LAN:	(config)#linkstate C-LAN
Verify configuration:	(config)#show linkstate

Private network

Private Network is a simple option of isolating IBP downlink ports (and thus servers) from each other. Isolated ports no longer have a connection to other downlink ports of the same group. Isolated ports can only establish a connection with IBP uplink ports/uplink-set.

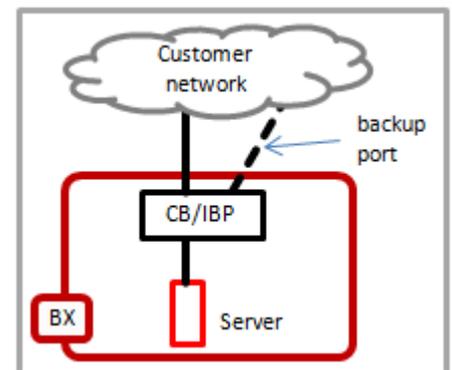
Configuration example

Create uplink-set C-LAN:	(config)#uplink-set C-LAN
Add IBP uplink ports to uplink-set C-LAN:	(Interface)#uplink-set C-LAN
Create port group Server and connect with uplink-set C-LAN:	(config)#port-group Server C-LAN
Remove IBP downlink port from default port group:	(Interface)#no port-group
Add IBP downlink ports to port group Server:	(Interface)#port-group Server
Configure IBP downlink port as private network:	(Interface)#isolate

Port backup

Port Backup enables for redundant connection of a single CB to two or more customer switches without the need for Spanning Tree Protocol. However, to achieve high availability it is recommended to at least use two IBPs.

The uplink port configured with port backup is the backup port for all ports of its uplink-set. That means, if an uplink connection fails (all of the physical connections, or logical, i.e. LAG/LACP), the backup port becomes active. It sends gratuitous arps of all active server mac-addresses to inform the upstream (customer) switch. Backup port assignment can be carried out with the uplink interface command **(interface)#port-backup**



Configuration examples

IBP default configuration

For non-stacked CBs, IBP provides a very useful configuration by default. With that, it is very easy to establish server connectivity without any VLAN administration. For stacked CBs, it is recommended to set up an individual IBP configuration.

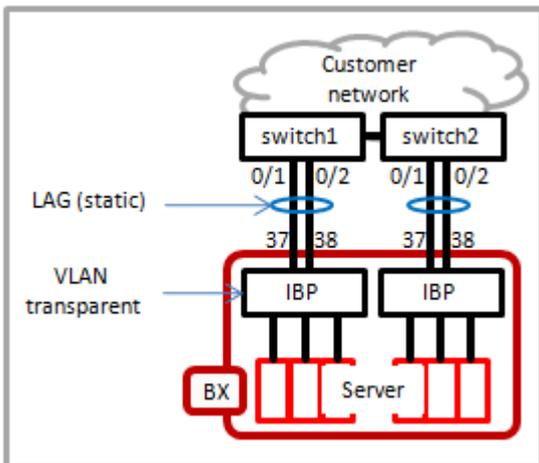
IBP default configuration:

- LAG (static)
The first eight RJ45 uplink ports (all CBs apart from CB18/8) respectively all uplink ports (CB18/8) form a LAG (static). That is because all these uplink ports are member of the uplink-set "default".
- VLAN transparency
VLANs are forwarded transparently; no VLAN administration is required. That means VLANs configured on the upstream (customer) switch are automatically allocated on all downlink (server) ports. That is because all downlink-ports are member of the port-group "default" which is associated with the uplink-set "default".
- Redundancy
Linkstate is enabled for the uplink-set "default" to inform the server teaming in case of an uplink failure.

Example 1: Individual IBPs with default configuration

In this example, no configuration is required on IBP site. On a newly delivered CB, just change the mode to IBP with the command #boot-system mode ibp. On customer switch site, just configure the LAGs (static) and required VLANs. All VLANs are automatically passed-thru to the blade servers respectively to the hypervisor virtual switch by the CB, for example VLANs 10-20 tagged / 1 untagged configured on the upstream switch are automatically allocated to all CB downlink (server) ports.

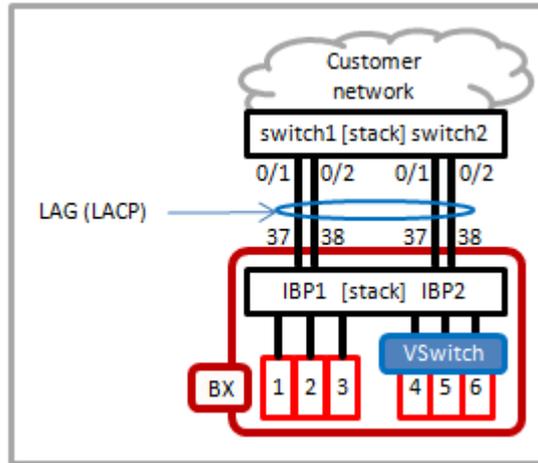
Example of upstream (customer) switch port configuration (cisco catalyst). The configuration is identical on both switches.



```
Switch#show running-config
!
Interface Port-channel1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 spanning-tree portfast trunk
!
Interface GigabitEthernet0/1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 spanning-tree portfast trunk
 channel-group 1 mode on
!
Interface GigabitEthernet0/2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 spanning-tree portfast trunk
 channel-group 1 mode on
```

Example 2: Stacked IBPs with mixed VLAN configuration

In this example, the CBs are stacked and connected via a single LAG (LACP) to the customer switch. All VLANs (in this case VLAN100-1000 tagged) configured on the customer switch port-channel are transparently forwarded to BX900 server 4, 5 and 6 hypervisor Vswitch; these VLANs can be administered on the VSwitch only, but do not need to be administered in the IBP. Server 1 is member of VLAN 101 (tagged), server 2 is member of VLAN 102 (tagged) and server 3 is member of VLAN 101 and 102 (tagged).



Example of upstream (customer) switch port configuration (cisco catalyst):

Detailed configuration steps of IBP:

```
Switch#show running-config
!
Interface Port-channel1
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100-1000
switchport mode trunk
spanning-tree portfast trunk
!
Interface GigabitEthernet1/0/1
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100-1000
switchport mode trunk
spanning-tree portfast trunk
channel-group 1 mode active
!
Interface GigabitEthernet1/0/2
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100-1000
switchport mode trunk
spanning-tree portfast trunk
channel-group 1 mode active
!
Interface GigabitEthernet2/0/1
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100-1000
switchport mode trunk
spanning-tree portfast trunk
channel-group 1 mode active
!
Interface GigabitEthernet2/0/2
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100-1000
switchport mode trunk
spanning-tree portfast trunk
channel-group 1 mode active
```

```
Create uplink-set C-LAN and activate LACP:
(config)#uplink-set C-LAN
(config)#lACP C-LAN
Add IBP uplink ports to uplink-set C-LAN:
(config)#interface range 1/0/37-1/0/38, 2/0/37-2/0/38
(if-range)#uplink-set C-LAN
(if-range)#exit
Create port-group and svc-vlan and associate each with uplink-set C-LAN:
(config)#port-group ESX C-LAN
(config)#svc-vlan PROD1 101 C-LAN
(config)#svc-vlan PROD2 102 C-LAN
Assign IBP downlink ports to downlink groups:
(config)#interface range 1/0/1, 2/0/1
(if-range)#no port-group
(if-range)#svc-vlan PROD1
(if-range)#exit
(config)#interface range 1/0/2, 2/0/2
(if-range)#no port-group
(if-range)#svc-vlan PROD2
(if-range)#exit
(config)#interface range 1/0/3, 2/0/3
(if-range)#no port-group
(if-range)#svc-vlan PROD1
(if-range)#svc-vlan PROD2
(if-range)#exit
(config)#interface range 1/0/4-1/0/6, 2/0/4-2/0/6
(if-range)#no port-group
(if-range)#port-group ESX
(if-range)#exit
(config)#
```

Restriction for CB18/8

For CB18/8 it is recommended to use Port-group and VLAN-group only. Also it is possible to use Service-LAN and Service-VLAN, these groups may degrade performance in some circumstances. Furthermore, Service-VLAN and Service-LAN by default have isolation-feature enabled, what means server-to-server communication is not possible within these groups. It is possible to disable the isolation-feature.

Mode independent information

CB file handling

Configuration files created with the **#copy** command ..

- optionally contains any configuration file
- are up-/downloadable to a ftp-/tftp-server with the **#copy** command
- are encrypted and are not readable / editable
- can be marked as "startup" configuration file

Configuration files created with the **#show run <filename.scr>** command ..

- are so called "scripted" files and contain the running-configuration only
- are up-/downloadable to a ftp-/tftp-server with the **#copy** command
- are readable and editable (not editable on local store)
- can't be marked as "startup" configuration file

#copy commands:

- To save the running-configuration to the startup-configuration (the local configuration file which has set its startup flag):

#copy running-config startup-config

Note: Startup flag is initially set to the file "default.cfg"

- To save the running-configuration to an individual configuration file and set the startup flag:

#copy running-config startup-config <filename>

Note: A maximum of 10 configuration files can be saved.

- To copy the startup-configuration (which is the local configuration file which has set its startup flag) to the running-configuration:

#copy startup-configuration running-configuration

Note: The CB / stack automatically reboots

- To set the startup flag of a configuration file: **#boot-system config <filename>**

- To show the files: **#dir**

- To delete a file: **#delete <filename>**

```
(C180cb2) #dir
```

date	file name	file type	startup	size (byte)
2009/04/16	sb11-b-1.00.0415.biz	Boot-Rom image	Y	373108
2011/11/17	sb11-sw-r-6.16.1117.biz	Operation Code	Y	8896069
2011/11/10	sb11-sw-r-6.14.1110.biz	Operation Code	N	8894047
2009/04/16	sb11-sw-factory-defaults	Config File	N	22124
2011/11/24	default.cfg	Config File	N	24512
2006/05/11	moore.cfg	Config File	N	26289
2010/01/20	bastian	Config File	N	24322
2011/11/21	no_chan.cfg	Config File	N	30072
2011/10/26	merged_stacks.cfg	Config File	N	29975
2011/10/26	ping_max.cfg	Config File	N	30087
2011/11/17	teams	Config File	N	30755
2011/11/21	matthias	Config File	Y	30072

Total: 12 files.

#show run <filename.scr> commands:

To save the running-config to an individual file: **#show run <filename>.scr**

To apply/delete/list/show files on local store: **#script ..**

```
(C180cb2) #script list
Configuration Script Name      Size(Bytes)
-----
TEST123.scr                    10963
TEST4711.scr                   3458

2 configuration script(s) found.
2033 Kbytes free.
```

Some more useful commands

To get all information of the CB (i.e. for support reasons): **#show tech-support**

To set the CB running-configuration to its defaults: **#clear config**

To set an interface to its default-configuration: **#clear config <interface>**

To get short interface names: **#classic view**

CB18/8**CB file handling**

The file system does not provide a command to view files. It is possible to store up to two individual configuration files (config1 and config2).

To copy files (configuration and image files for update): **#copy**

```
c180cb1#
c180cb1# copy
  <src_filename>                Specify a source file name
  <tftp://[ipv6address]/filename> Download the file by TFTP (IPv6)
  <tftp://ipv4address/filename>  Download the file by TFTP (IPv4)
  config1                       Configuration file 1
  config2                       Configuration file 2
  ibp_firm                      IBP firmware
  running-config                Configuration during operation
  startup-config                Configuration to be used for device
                                startup
  switch_firm                   Switch firmware
c180cb1# copy
```

To view all system information: **#show system information**

```
c180cb1#
c180cb1# show system information
Current-time : Sun Apr 22 11:23:02 2012
Startup-time : Fri Apr 20 16:37:57 2012
Hardware Version : CO-07
Number of ports : 18/8
Machine Model : PY CB Eth Switch/IBP 10Gb 18/8
ROM Ver. : 2.4
ASIC Firm Ver. : 090814PL21
Firm Ver. : V03.00 NY0061 Mon Apr 9 13:38:11 JST 2012
Startup-firmware : ehm_firm
Next-boot : switch_firm
Startup-config : Thu Jan 1 00:00:00 1970 config1
Running-config : Sun Apr 22 10:38:22 2012
MAC : 001742d113e0
Memory : 128MB
c180cb1#
```

Some more useful commands

To get all information of the CB (i.e. for support reasons): **#show tech-support**

To set the CB configuration to its defaults: **#reset clear**

To display a difference of configuration files: **#diff**

To set Ethernet ports offline (link status "down"): **#offline**

CB18/6, CB36/xx

Convergence time after a stack-master failure

With stacked CBs (stacking is possible with CB36/8+2 only), if there is a CB stack-master failure, there is a convergence time in which the whole stack does not forward traffic and drops it. The convergence time depends primarily on the activated Mode (EHM|IBP|Switch) and on the amount of stack member.

# of stack member	EHM	IBP-Mode	Switch-Mode
2	< 1 second	< 3 seconds	< 1 second
4	< 10 seconds	< 10 seconds	< 1 second
> 4	< 20 seconds	< 20 seconds	< 10 seconds

Overview of ETH Connection Blades supporting IBP

It is possible to switch the following Ethernet CBs into one of the three modes EHM/IBP/SWICH via software command, whereas EHM is the factory default mode. All CBs fit into the BX400 and BX900 chassis.

PY CB Eth Switch/IBP 1Gb 18/6 (codename „SB6“)

18 x 1GE downlink
6 x 1GE uplink (RJ45)



PY CB Eth Switch/IBP 1Gb 36/12 (codename „SB11a“)

36 x 1GE downlink
12 x 1GE uplink
Ports 37-40 (requires SFP modules)
Ports 41-48 (RJ45)



PY CB Eth Switch/IBP 1/10Gb 36/8+2 (codename „SB11“)

36 x 1GE downlink
10 x uplink
2 x 10GE: Ports 37-38 (requires SFP+ modules)
8 x 1GE: Ports 39-46 (RJ45)
1 x stacking CX4-Port (external)
1 x stacking-Port (internal)



PY CB Eth Switch/IBP 10Gb 18/8 (codename „SBAX2“)

18 x 1/10GE downlink
8 x 1/10GE uplink (requires SFP/SFP+ -Modules)



Changing the CB Mode

EHM is the default mode for all CBs delivered after April 2012 (CB18/8: June 2012). Before that date, switch-mode is the default mode.

To change the CB mode `#boot-system mode {EHM|IBP|switch}`

To verify the mode(CB18/6, CB36/xx)`#whichboot (CB18/8)#show system information`

BX400/BX900 midplane connectivity

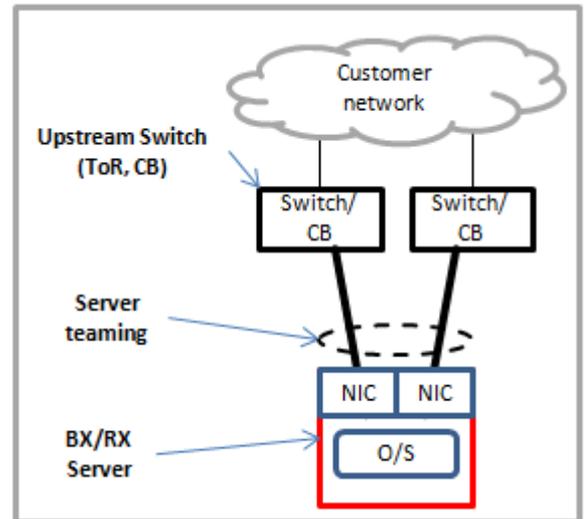
Connections between Blade-Servers and Connection Blades are based on the BX400/BX900 midplane. To get an overview of the physical connections between server adapter and Connection Blade, an Excel-sheet is provided. You can fill in the vmnic (if applicable) to get an end-to-end view within the BX400/BX900 chassis.

[Link to BX400/BX900 midplane connectivity excel sheet](#)

Server teaming options

Server teaming configuration depends on several factors, these are...

- Server O/S (Operating System).
Linux and VMware use its own teaming driver, while Windows uses the adapter teaming driver (Broadcom, Intel).
- Server Ethernet adapter (NIC "Network Interface Card").
Both Broadcom and Intel adapter offers in general similar teaming functions but differ in the naming of the functions.
- Upstream switch.
Might be the BX400/BX900 ETH CB, or in case of RX server the ToR (Top of Rack) switch which is usually the customer switch. It has to be considered whether the upstream switches are standalone or clustered (stacking, Brocade MLAG/vLAG, Cisco vPC/VSS, etc...).

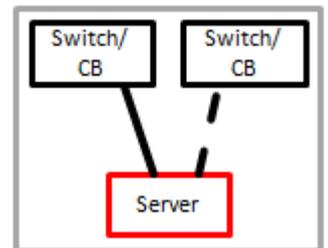


There are four different teaming options available:

- Active/standby
In an active/standby configuration, the standby path is the redundancy for the active path; it does not send traffic and ignores received traffic. No dedicated configuration is required on the upstream switch. A physical connection to different upstream switches is possible.

In case of teaming is provided by the adapter driver (i.e. Windows, except Server 2012), the configuration can be carried out with the following option:

- Intel: Adapter Fault Tolerance
- Broadcom: Smart Load Balancing



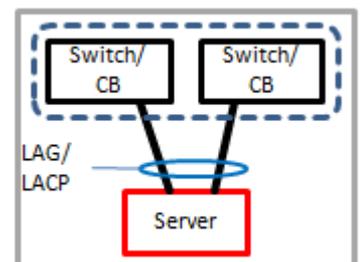
In case of teaming is realized with the O/S driver, the configuration can be carried out with the following option:

- Linux: active-backup (1)
- VMware: Explicit Failover Order
- Cisco Nexus1000v: not supported

- Active/active with LAG (Link Aggregation Group) standard 802.3ad
On an 802.3ad standard LAG, LACP (Link Aggregation Control Protocol) is enabled. Up to 8 links can be active at the same time. Dedicated configuration is required on the upstream switch (Cisco Interface command `#channel-group xy mode active`). A physical connection to a single upstream switch (cluster) is possible only.

In case of teaming is provided by adapter driver (i.e. Windows, except Server 2012), the configuration can be carried out with the following option:

- Intel: Dynamic Link Aggregation
- Broadcom: Link Aggregation (802.3ad)

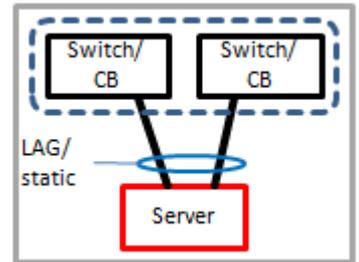


In case of teaming is realized with the O/S driver, the configuration can be carried out with the following option:

- Linux: 802.3ad (4)
- VMware: supported with v5.1 dVS
- Cisco Nexus1000v: „channel-group x mode active“

- Active/active with LAG (static)

Also called Etherchannel, a static LAG is similar to a standard 802.3ad LAG but does not have any protocol enabled between the switches (like LACP). Dedicated configuration is required on the upstream switch (Cisco Interface command `#channel-group xy mode on`). A physical connection to a single upstream switch (cluster) is possible only.



In case of teaming is provided by adapter driver (i.e. Windows, except Server 2012), the configuration can be carried out with the following option:

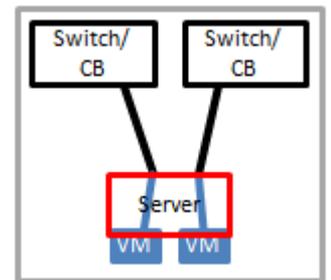
- Intel: Static Link Aggregation
- Broadcom: Generic Trunking (FEC/GEC) 802.3ad-Draft Static

In case of teaming is realized with the O/S driver, the configuration can be carried out with the following option:

- Linux: not supported
- VMWare: IP Hash
- Cisco Nexus1000v: „channel-group x mode on“

- Active/active without LAG

With that option, VMs are pinned to dedicated Server NICs and appear to the upstream switch / CB as dedicated, directly connected servers. Although this option is available for physical server, FTS in general recommends using it for virtualized servers only.



In case of teaming is provided by adapter driver (i.e. Windows, except Server 2012), the configuration can be carried out with the following option:

- Intel: Adaptive Load Balancing/Receive Load Balancing (ALB/RLB)
- Broadcom: Smart Load Balancing

In case of teaming is realized with the O/S driver, the configuration can be carried out with the following option:

- Linux: balance-alb (6)
- VMware: „Originating Virtual Port ID“ (recommended) or „Source MAC Address“
- Cisco Nexus1000v: Mac pinning

Contact

FUJITSU LIMITED
Website: www.fujitsu.com/fts
2014-10-10 WW-EN

© Copyright 2014 Fujitsu Limited. Fujitsu, the Fujitsu logo are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.